

WHAT IS CLAIMED IS:

1. A scalable clustered system, comprising:

a global fabric; and

two or more cluster nodes interconnected via the global fabric, each cluster node including a node naming agent (NNA), a local fabric and one or more end nodes interconnected via the local fabric, the NNA being configured as a fully symmetrical translation device interposed between the local fabric and the global fabric, the NNA providing support for scaled clustering by transforming a local cluster address into a corresponding global cluster address for each packet in an outbound path from any of the cluster nodes and by transforming a global cluster address into a corresponding local cluster address for each packet in an inbound path to any of the cluster nodes, wherein intra-node cluster addressing is transparent to inter-node cluster address changes, and wherein re-configuration of the scalable clustered system requires no address re-assignments yet allowing the end nodes in the cluster nodes to maintain connectivity therebetween.

2. A scalable clustered system as in claim 1 wherein the local fabric and global fabric provide local and global clustering support infrastructures, respectively, and wherein global routing and global fabric topology are transparent to end nodes in operation.

3. A scalable clustered system as in claim 1 wherein the local and global fabric are each configured with one or more routers and/or switches.

4. A scalable clustered system as in claim 1 wherein each of the end nodes is an addressable device representing a resource such as a microprocessor, a central processing unit (CPU), a memory, an input/output (I/O) device controller or a server.

5. A scalable clustered system as in claim 1, the re-configuration of which, by addition thereto or removal therefrom of any cluster node, is possible while maintaining the configuration of any of the other cluster nodes, and wherein configuration of any of the cluster nodes is possible without reconfiguration of any of the other cluster nodes.

6. A scalable clustered system as in claim 1 wherein connections in the local and global fabric can be modified while maintaining correct packet transmission behavior.

7. A scalable clustered system as in claim 1 wherein the number of end nodes in each cluster node corresponds to an attribute of address fields in packets.

8. A scalable clustered system as in claim 1 wherein the NNA is configured to prevent a duplicated assignment of any end node address in establishing end node membership in a cluster node.

9. A scalable clustered system as in claim 1 wherein the scalable clustered system has a hierarchical topology and wherein each of the cluster nodes has either a flat or a hierarchical topology.

10. A scalable clustered system as in claim 1 wherein local traffic of packets within any of the cluster nodes is not routed to its respective NNA.

11. A scalable clustered system as in claim 1 wherein global traffic of packets to and from each of the cluster nodes is routed via its respective NNA.

12. A scalable clustered system as in claim 1 wherein the NNA includes a mask register for transforming the global/local cluster addresses in which bit substitutions can be made before an entire address has arrived at the NNA.

13. A scalable clustered system as in claim 1 having hierarchical topology and address-identification scheme that relative to flat non-hierarchical topology require smaller address fields and routing tables.

14. A scalable clustered system as in claim 1 wherein the NNA is a semiconductor chip.

15. A scalable clustered system as in claim 1 wherein the NNA has symmetrically built ends with one end being connected to the local fabric and the other end being connected to the global fabric and wherein a capability replicated at each end includes status indications and control enable bits.

5

16. A scalable clustered system as in claim 15 wherein the control enable bits include a shutdown on missing clock enable, a replace source address enable, a replace destination address enable, a destination address checking enable and a pass-through enable.

10

17. A scalable clustered system as in claim 1 wherein the NNA has a mode control register the contents of which determines which mode of operation the NNA assumes.

15
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95
100

18. A scalable clustered system as in claim 17 wherein the NNA modes of operation include a pass-through mode, a conversion mode, a clock-check mode, an error check mode an error recovery mode and a shutdown mode.

19. A scalable clustered system as in claim 1 wherein, while in the conversion mode, the NNA is configured to perform source or destination cluster identification (ID) translation, clock checking and destination address checking.

20. A scalable clustered system as in claim 18 wherein the pass-through mode is a default mode of the NNA upon initialization.

25

21. A scalable clustered system as in claim 1 wherein NNA includes data replacement registers programmable with information for converting local cluster address to global cluster address and global cluster address to local cluster address.

30

22. A method for scaling a clustered system, comprising:
operatively linking two or more cluster nodes via a global fabric in order to form a larger clustered system, each of the cluster nodes having end nodes and a local fabric interconnecting the end nodes;

routing global packet traffic between the two or more cluster nodes in the larger clustered system via the global fabric;

routing local packet traffic between the one or more end nodes within each of the cluster nodes via the local fabric; and

5 operatively interposing an NNA between the local fabric and the global fabric, the NNA being configured as a fully symmetrical translation device, the NNA providing support for scaled clustering by transforming a local cluster identification (ID) into a corresponding global cluster ID for each packet in an outbound path from any of the cluster nodes and by transforming a global cluster ID into a corresponding local cluster ID for each packet in an inbound path to any
10 of the cluster nodes, wherein intra-node cluster addressing is transparent to inter-node cluster address changes, and wherein re-configuration of the scalable clustered system requires no ID re-assignments yet allowing the end nodes in the cluster nodes to maintain connectivity therebetween.

15 23. A method as in claim 22 wherein in the outbound path the NNA performs steps including forwarding control and status symbols,
 replacing a local cluster ID value with a global cluster ID value in a packet source address field if the NNA operates in a conversion mode, and
 establishing a cyclic redundancy check (CRC) value in an outbound packet CRC
20 field based on whether a correct or incorrect CRC value is detected.

25 24. A method as in claim 23 wherein the CRC value is re-computed if a correct CRC value is detected, a destination cluster ID is enabled and there is no destination cluster ID mismatch.

30 25. A method as in claim 24 wherein in the inbound path, the NNA performs steps including forwarding control and status symbols,
 verifying proper routing by checking a destination cluster ID field in an inbound packet,
 detecting a correct or incorrect CRC value in the inbound packet,
 setting a status bit if the incorrect CRC value is detected, and

replacing a global cluster ID value with a local cluster ID value in a packet destination address field if the NNA operates in a conversion mode.

26. In a scalable clustered system, a computer readable medium embodying computer program code configured to cause that system to perform steps for configuring and scaling that system, comprising:

operatively linking two or more cluster nodes via a global fabric in order to form a larger clustered system, each of the cluster nodes having end nodes and a local fabric interconnecting the end nodes;

routing global packet traffic between the two or more cluster nodes in the larger clustered system via the global fabric;

routing local packet traffic between the one or more end nodes within each of the cluster nodes via the local fabric; and

operatively interposing an NNA between the local fabric and the global fabric, the NNA being configured as a fully symmetrical translation device, the NNA providing support for scaled clustering by transforming a local cluster address into a corresponding global cluster address for each packet in an outbound path from any of the cluster nodes and by transforming a global cluster address into a corresponding local cluster address for each packet in an inbound path to any of the cluster nodes, wherein intra-node cluster addressing is transparent to inter-node cluster address changes, and wherein re-configuration of the scalable clustered system requires no address re-assignments yet allowing the end nodes in the cluster nodes to maintain connectivity therebetween.

27. A computer readable medium as in claim 26 wherein the computer program code is further configured to cause the NNA to prevent

assignment of a same local address to two end nodes in the same cluster node, and associating an end node with more than one cluster node.

28. A scalable clustered system, comprising:

global interconnection means for operatively linking two or more cluster nodes in order to form a larger clustered system, each of the cluster nodes having end nodes;

local interconnection means for operatively linking the end nodes within each of the cluster nodes;

global routing means for routing global packet traffic between the two or more cluster nodes in the larger clustered system via the global interconnection means;

5 local routing means for routing local packet traffic between the one or more end nodes within each of the cluster nodes via the local fabric means; and

translation means operatively interposed between the local interconnection means and the global interconnection means, the translation means being configured as a fully symmetrical translation device, the translation means providing support for scaled clustering by transforming
10 a local cluster address into a corresponding global cluster address for each packet in an outbound path from any of the cluster nodes and by transforming a global cluster address into a corresponding local cluster address for each packet in an inbound path to any of the cluster nodes, wherein intra-node cluster addressing is transparent to inter-node cluster address changes, and wherein re-configuration of the scalable clustered system requires no address re-assignments
15 yet allowing the end nodes in the cluster nodes to maintain connectivity therebetween.

29. A scalable super-clustered system, comprising:

a primary level of hierarchy configured with a plurality of global fabrics each of which interconnecting a plurality of cluster nodes to form one or more primary-level clusters, each cluster node including a node-level node naming agent (NNA), a local fabric and one or more
20 end nodes interconnected via the local fabric, the node-level NNA being configured as a fully symmetrical translation device interposed between its local fabric and one of the global fabrics to which the cluster node is connected; and

an upper level of the hierarchy configured with one or more upper-level global fabrics
25 each of which interconnecting a plurality of the primary-level clusters to form one or more upper-level clusters, each primary-level cluster including a primary-level NNA that is configured as a fully symmetrical translation device interposed between its global fabric and a particular one of the upper-level global fabrics to which the primary-level cluster is connected,

wherein intra-node cluster addressing is transparent to inter-node cluster address changes
30 at any level of the hierarchy, and wherein re-configuration of the scalable super-clustered system

requires no address re-assignments yet allowing the end nodes in the cluster nodes to maintain connectivity therebetween.

30. A scalable super-clustered system as in claim 29, wherein the node-level NNA is configured to provide support for super-scaled clustering by transforming a local primary-level cluster address into a corresponding upper-level global cluster address for each packet in an outbound path from any of the primary-level clusters and by transforming an upper-level global cluster address into a corresponding local primary-level cluster address for each packet in an inbound path to any of the primary-level clusters.

31. A scalable super-clustered system as in claim 29, wherein the primary-level NNA is configured to provide support for super-scaled clustering by transforming a local cluster address into a corresponding global cluster address for each packet in an outbound path from any of the cluster nodes and by transforming a global cluster address into a corresponding local cluster address for each packet in an inbound path to any of the cluster nodes.

32. A scalable super-clustered system as in claim 29, wherein the NNA is configured with a software algorithm to prevent a duplicate assignment of any address so that, at any point in time, an end node cannot be a member of more than one of the plurality of cluster nodes.

33. A scalable super-clustered system as in claim 29 capable of operating in a TCP-IP (transmission control protocol – Internet protocol) network environment.

34. A scalable clustered system as in claim 1, wherein the NNA is configured with a software algorithm to prevent a duplicate assignment of any address so that, at any point in time, an end node cannot be a member of more than one of the plurality of cluster nodes.

35. A scalable clustered system as in claim 1 capable of operating in a TCP-IP network environment.

36. A scalable clustered system as in claim 28 capable of operating in a TCP-IP (transmission control protocol – internet protocol) network environment.

37. A method as in claim 22 implemented in a TCP-IP (transmission control protocol – internet protocol) network environment.

38. A scalable super-clustered system as in claim 29 further comprising:
higher levels of the hierarchy each of which including a higher-level fabric interconnecting clusters of the level below, each such cluster including a sub-cluster NNA interposed between that cluster and the higher-level fabric to which it is connected.

39. A scalable super-clustered system as in claim 29, wherein the local cluster address is a fixed cluster number of a cluster node, and wherein the global cluster address is a cluster number assigned to the cluster node during cluster configuration.

40. A scalable clustered system as in claim 1, wherein the local cluster address is a fixed cluster number of a cluster node, and wherein the global cluster address is a cluster number assigned to the cluster node during cluster configuration.